

WHY MACHINES CAN'T THINK:
A REPLY TO JAMES MOOR

(Received 27 September, 1977)

In 'An Analysis of the Turing Test,' James Moor claims that all too many have simply misunderstood how Turing's test figures in arguments about the mentality of machines.¹ The test is a familiar one: an interrogator enters Turing's question/answer setup with the aim of finding out which respondent is another person, which a computer. According to Moor, this sort of test can be "interpreted inductively" (p. 266). That is, one can perfectly well view Turing's test as providing "behavioral evidence" (p. 253). Indeed, Moor thinks it can provide enough evidence to secure the point at issue. He takes the test results, when passing, as evidence of an ability to think. Passing, in Moor's sense, comes to this: many different interrogators are allowed as many chances to question as they like, and yet in the end an average interrogator can only spot the machine about 50% of the time (pp. 249–50). Moor thinks this would be "very adequate grounds for inductively inferring that the computer could think" (p. 251). Would it be?

Though Moor calls his interpretation an inductive one, it is really more accurate to call it an explanatory one. This becomes clear when he discusses why we should take a computer's behavior as telling evidence of cognition. Moor first turns to one's own situation with respect to other people. He wants us to pay attention to something most believe: that other people can and do think. Why do we believe this? Moor finds his answer by appealing to a theory:

I believe that another human being thinks because his ability to think is part of a theory I have to explain his actions. The theory postulates a number of inner information processes, but the evidence for the theory comes from the outward behavior of the person. On the basis of his behavior I can confirm, disconfirm, and modify my theory. (p. 251)

On this approach, one's beliefs about the mentality of others are part of an explanatory theory. In order to explain the behavior of others, we invoke a theory that involves the notion of thinking. But this isn't, so far, the full

story. It doesn't tell us why we should take a person's behavior as telling evidence of a certain mental life. Going on this explanatory approach, the behavior counts as evidence because it is connected with a going theory. How does it count as telling evidence? To be that, it needs to be connected with the best of the going theories. When it comes to everydayish efforts at explaining the behavior of other people, one is hard put to find anything better than the current mentalistic scheme. That scheme involves, of course, one's common notion of thinking. And it's presumably the scheme that Moor relies on here.

When Moor turns to computers, he urges a parity. He claims that our situation with respect to other people is the same as ours with respect to computers. We need to explain the computer's behavior, and so we invoke a theory. As Moor puts it:

Furthermore, there is no reason why knowledge of computer thinking can not arise in the same way. I can use the computer's behavior as evidence in assessing my theory about its information processing. (p. 251)

Moreover, the computer's behavior is the same kind of behavior we take as evidence that other people can think: "the Turing test permits direct or indirect testing of virtually all of the activities one would count as evidence for thinking" (p. 251). For example, it provides a direct way to check on a computer's verbal behavior. With its question/answer format, Turing's test "permits (even demands) evaluation of linguistic behavior which is central to our inductive inferences about how others think" (p. 251). It also provides an indirect way to check on nonverbal behavior. An interrogator can ask for descriptions of how the respondent would do something that takes some thinking (p. 252). As Moor sees it, what counts as telling evidence for people also counts as that for computers. Thus he invokes a theory that involves the notion of thinking; and this supposedly explains the behavior of a computer that can pass Turing's test.

Let's grant that an explanatory approach is a viable one for questions of computer cognition. Even so, Moor arrives at his conclusion all too quickly. He glosses over a step that one simply can't pass by. To put the point another way, Moor leaves an essential step unargued and assumed: viz., that his theory for explaining a computer's behavior is better than others about. Without that step, Moor's argument is really no argument at all. For example, Moor takes a computer's linguistic behavior as evidence in the way that a person's linguistic behavior is evidence. He counts both bits of behavior as decided evidence of some thinking on the part of each. It is, Moor claims, just what "one would